arXiv:1609.00851v1 [cond-mat.stat-mech] 3 Sep 2016

Inverse Design for Self Assembly via On-the-Fly Optimization

Beth A. Lindquist,[1] Ryan B. Jadrich,[1] and Thomas M. Truskett[1, a)]
McKetta Department of Chemical Engineering, University of Texas at Austin, Austin, Texas 78712, USA

(Dated: 6 September 2016)

Inverse methods of statistical mechanics have facilitated the discovery of pair potentials that stabilize a wide variety of targeted lattices at zero temperature. However, such methods are complicated by the need to compare, within the optimization framework, the energy of the desired lattice to all possibly relevant competing structures, which are not generally known in advance. Furthermore, ground-state stability does not guarantee that the target will readily assemble from the fluid upon cooling from higher temperature. Here, we introduce a molecular dynamics simulation-based, optimization design strategy that iteratively and systematically refines the pair interaction according to the fluid and crystalline structural ensembles encountered during the assembly process. We successfully apply this probabilistic, machine-learning approach to the design of repulsive, isotropic pair potentials that assemble into honeycomb, kagome, square, rectangular, truncated square and truncated hexagonal lattices.

There is a growing appreciation that a diverse array of structural motifs can be stabilized in systems of particles interacting via isotropic pair potentials, including various microphases[1–3], open crystalline lattices[4–7], and quasi-crystals[8,9]. Discovery of such potentials has been facilitated by inverse methods of statistical mechanics, most commonly optimization algorithms that iteratively refine the form of the interaction to attain increasingly favorable structural or thermodynamic properties.[10,11] For cases where such strategies have been employed to find isotropic pair potentials that assemble into specific two- and three-dimensional (2D and 3D) lattices, optimization typically builds on a ground-state calculation, wherein interactions are sought that stabilize the target structure at zero temperature relative to relevant competing structures. Despite successful application of such methods to design purely repulsive interactions stabilizing multiple open lattices (including honeycomb,[5–7] square,[5–7] rectangular,[4] and kagome[4] in 2D and diamond,[5,12,13] simple cubic,[5,12] and fluorite[4] in 3D), there are some notable drawbacks to these approaches. First, they are encumbered by the requirement to specify the pool of relevant competing structures, a list that is not fully known in advance and thus must be modified as the potential is updated in the optimization. Moreover, the target structure must be checked explicitly for mechanical stability with the pair potential. Finally, interactions designed to stabilize the target structure in the ground state are not guaranteed to readily assemble the target from the fluid state upon cooling from higher temperature.

In this Communication, we report a molecular dynamics (MD) simulation-based, inverse optimization strategy–carried out at nonzero temperature–that iteratively and systematically refines the pair interaction according to the fluid and crystalline pair structures dynamically encountered during the assembly process. The approach, while encoding practical aspects of assembly of the target from the fluid, is also technically simple and easy to implement. To illustrate the power of the methodology, we successfully employ it to design isotropic and purely repulsive pair potentials to assemble six distinct 2D lattices, including two structures which–to our knowledge–have never been stabilized via a pair potential before.

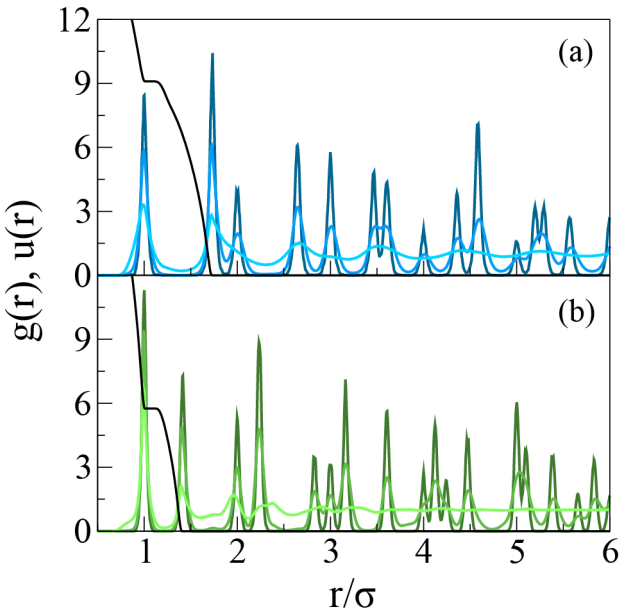

FIG. 1. For the HC (a) and SQ (b) lattices, the optimized potential, $u(r)$ (thin black line), and the radial distribution functions, $g(r)$, corresponding to (from strongest to weakest structuring in the first coordination shell): the target simulation, simulation with the optimized potential, and the fluid in the final step in the optimization prior to crystallization.

Machine learning-based techniques have been successfully applied to the self-assembly of colloids, from the rational design of building blocks and templates needed to fabricate nanomaterials[15–17] to the elucidation of pathways involved in the assembly process.[18] In this work, the

a)Electronic mail: truskett@che.utexas.edu

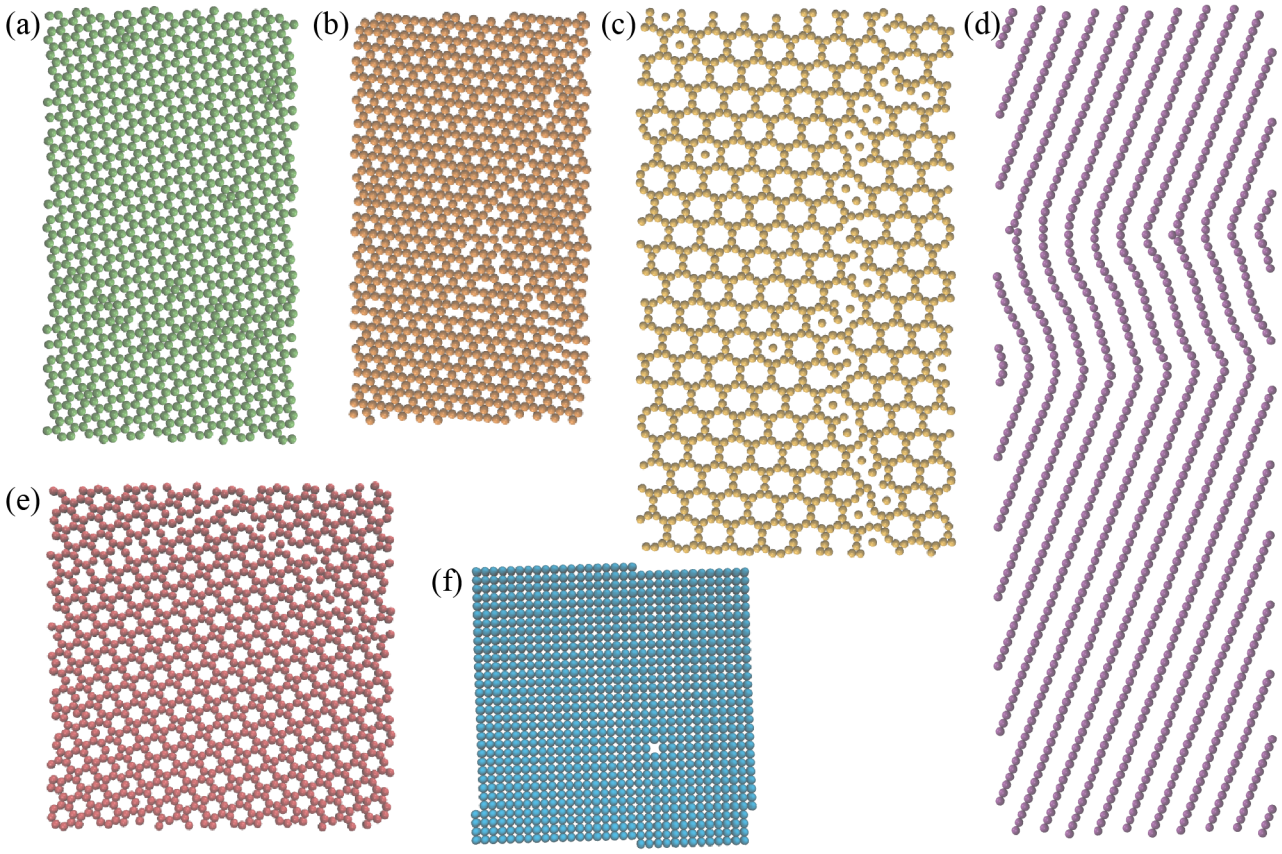

FIG. 2. Final configurations after cooling to $T = 0$ with $u(r)$ designed for the (a) honeycomb (HC), (b) kagome (KG), (c) truncated hexagonal (TH), (d) rectangular with an aspect ratio of three (R3), (e) truncated square (TS), and (f) square (SQ) lattices. All configurations were visualized in VMD.[14]

optimization scheme we present is general and based on maximum-likelihood machine learning[19] (called 'relative entropy coarse-graining' in the statistical mechanics community[20–22]). Within it, particle-particle interactions are tuned in order to maximize the likelihood of reproducing desired configurations (in this specific case, 2D periodic lattices). We have previously used a related optimization strategy based upon relative entropy coarse-graining to discover isotropic pair potentials that favor microphases with porous architectures of a prescribed size.[3]

One advantage of relative entropy coarse-graining is the expression of the pair potential $u(r|\boldsymbol{\theta})$) in terms of a functional form that is parametrized by m arbitrary, scalar parameters, $\boldsymbol{\theta} \equiv [\theta_1, \theta_2, ..., \theta_m]$; this allows for a variety of constraints to be straightforwardly placed on $u(r|\boldsymbol{\theta})$. For flexibility of the potential, we optimize the amplitudes of the knots for an Akima spline spaced at an interval (δr) of $\delta r/\sigma \approx 0.0056$ where σ is the nearest neighbor crystal distance. In order to avoid unphysical oscillatory potentials,[23] we constrain the knots to be monotonically increasing with decreasing r (i.e., purely repulsive interactions), though other more complex functional forms and constraints are possible. For the case of isotropic pair potentials, the updates to the parameters that characterize the potential are derived from the difference in the radial distribution functions associated with the present potential $g(r|\boldsymbol{\theta})$ and the target simulation, $g_{\mathrm{tgt}}(r)$:

$$\boldsymbol{\theta}^{(i+1)} = \boldsymbol{\theta}^{(i)} + \alpha \int_0^\infty dr r \left[g(r|\boldsymbol{\theta}^{(i)}) - g_{tgt}(r)\right] \left[\boldsymbol{\nabla}_{\boldsymbol{\theta}} u(r|\boldsymbol{\theta})\right]_{\boldsymbol{\theta}=\boldsymbol{\theta}^{(i)}} \quad (1)$$

where i indexes the iteration and α is the learning rate to be set empirically by observing the optimization stability.[24] A derivation of the update scheme and additional details pertaining to the optimization algorithm can be found in the Appendix.

The target configurations are generated by an MD simulation where the particles are pinned to their lattice positions via a quadratic confining potential;[25] all simulations contained ≥ 1000 particles and were performed in Gromacs 4.6.5[26,27] in the NVT ensemble[28] with periodic boundary conditions in the x and y directions. For notational convenience, we define our optimization temperature as T^*; however, the actual outcome of the optimization is $\beta u(r)$, which uniquely defines the $u(r)$ at any given temperature ($\beta = (k_B T)^{-1}$) that yields the desired

pair structure. For the analysis below, the $u(r)$ given by T^* is used, where k_BT^* assumes a value of unity. Further details regarding simulations can be found in the Appendix.

In contrast to prior work based on finding a potential for which the desired lattice is the ground state, we target potentials that self-assemble into the targeted lattice at finite temperature. Therefore, a key step in our procedure is to initiate each simulation from a disordered fluid state, thereby allowing all relevant structural motifs to compete as needed *prior to crystallization*, including structures such as disordered microphases that might not be easily amenable to inclusion in an explicit competitor pool. Moreover, mechanical stability is directly incorporated into the optimization framework by the presence of finite temperature thermal motion. Practically, the configuration from the previous step in the optimization is melted via heating to an empirically determined temperature of $T/T^* = 1.5$, and then subsequently cooled to $T/T^* = 1.0$ prior to collecting statistics for the $g(r)$. We consider the optimization complete when the crystal is sufficiently stable that it no longer melts at $T/T^* = 1.5$.

In Fig. 1a,b, we compare the radial distribution functions of the optimized ($g(r|\boldsymbol{\theta})$) and target ($g_{\rm tgt}(r)$) structures for both the honeycomb (HC) and square (SQ) lattice, respectively. Both show excellent matching in the peak positions over many coordination shells, even though the ranges of the potentials ($u(r)$, in black) only span the first two coordination shells. We also show $g(r)$ from the last step before any crystallization occurred in the optimization to demonstrate that the disordered fluid locally contains muted structural signatures of the lattice. This correspondence between disordered fluid and crystalline lattice structure allows for the former to provide *implicit* "competitor pool" information for the optimization. As the potential is optimized prior to crystallization and therefore the fluid state evolves, fluids that have structural signatures corresponding to other lattices (and therefore do not match $g_{\rm tgt}(r)$) are suppressed (any incorrectly assembled structures are also explicitly penalized). However, because the fluid is globally disordered and highly mobile, issues such as phase boundaries and metastability that become problematic upon lattice formation do not inhibit proper sampling of phase space in the fluid.

In addition to HC and SQ, this optimization procedure was carried out for the rectangular lattice with an aspect ratio of three (R3), kagome (KG), truncated square (TS), and truncated hexagonal (TH) lattices. The resulting potentials were then simulated shortly at a sufficiently high temperature to fully melt the crystal, and then the simulation was very slowly cooled through the empirically determined melting-freezing transition range and then further cooled to $T = 0$. The resulting structures are shown in Fig 2a-f, where it is clear that the optimizations were successful. Only small expected defects are present due to finite size of the periodically replicated simulation cell and any misalignment of the nucleated crystal with the cell. The corresponding potentials are shown in Fig. 3. The potentials are essentially repulsive shoulders, with one to three such features beyond the core. Shoulders originate from the monotonicity constraint, which prevents the development of attractive wells at specific coordination shells. Instead, $u(r)$ develops features with stiff repulsive forces 1) to penalize the surmounting of the shoulder, while 2) yielding a strong "thermodynamic" pressure (i.e., the ensemble averaged force due to the influence of every other particle) that pushes two particles into the shoulder. In essence, attractive wells are forgone for stiff repulsions and higher pressures to build the strong, specific coordination needed for targeting a crystal phase.

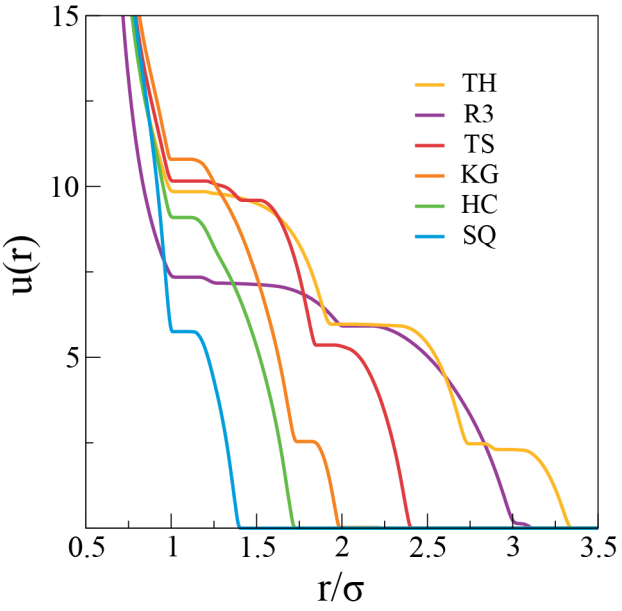

FIG. 3. Optimized pair potentials for the lattices considered in this work. The figure legend is ordered by decreasing range in $u(r)$ from top to bottom.

While the optimization framework presented here does not guarantee that the assembled crystal is also the ground state, we can confirm that the desired lattice, where the particles are in their ideal positions, is the lowest energy state at $T = 0$ of the lattices featured in this work. In Fig. 4a,b, we show the total energy, U, as a function of ρ for the six lattices studied here in addition to the triangular (TR) lattice, using $u_{\rm HC}(r)$ and $u_{\rm SQ}(r)$, respectively. The circle denotes the optimization density for the lattice, and we see that the desired lattice is the lowest in energy for a reasonable range about ρ in the optimization, though the values of ρ outside of the optimization point are not guaranteed to be mechanically stable. The other potentials display similar behavior (data not shown).

The importance of heating the system into the fluid phase at each optimization step, thereby incorporating the assembly process into the optimization, can be demonstrated in the context of the TH lattice. If opti-

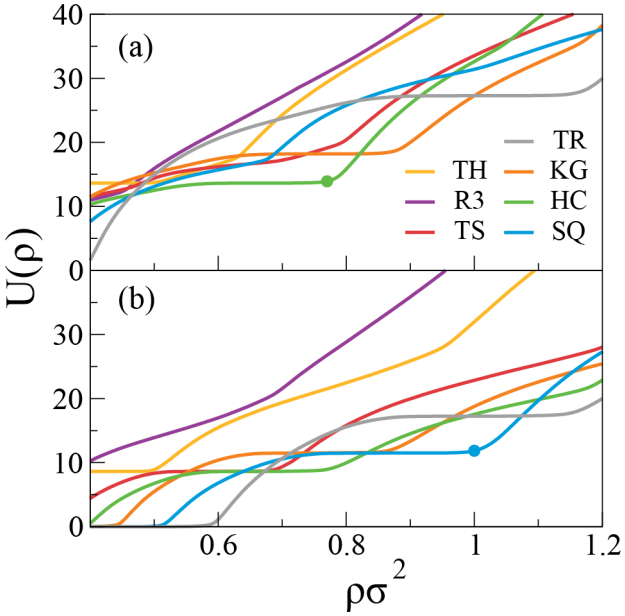

FIG. 4. Ground state energies for the lattices optimized in this work, plus a triangular lattice, for (a) $u_{\mathrm{HC}}(r)$ and (b) $u_{\mathrm{SQ}}(r)$ as a function of ρ. The circle denotes the density where the optimization was performed, ρ_{opt}. At ρ_{opt}, the potentials are ordered as (a) HC, KG, TS, SQ, TR, TH, R3 and (b) SQ, TR, HC, KG, TS, TH, R3 with increasing energy.

mizations are performed beginning from the crystalline state, then mechanical stability is accounted for since the crystal can fall apart in the MD simulation, but the role of competitor phases to the self-assembly process *is not included*. Simulated annealing of a potential optimized for the TH lattice without melting the crystal at the outset of each iteration resulted in assembly of a crystalline stripe phase as shown in Fig. 5. The energies at $T = 0$ of the resulting (defective) striped phase and the perfect TH lattice are similar, with the latter being more stable by 5%. However, the striped phase forms first and is sufficiently kinetically stable that it never transitions into the TH lattice, even with very slow annealing schedules. Therefore, the TH phase appears to be kinetically inaccessible via simulation, though it might in fact be the ground state. Beginning with a disordered state in every optimization step circumvents this difficulty entirely because the $g(r)$ is collected from the assembled phase, and the update scheme therefore adjusts the potential accordingly if an incorrect structure is encountered.

In summary, we have introduced a simple, relative entropy based inverse design approach, which we demonstrated can successfully discover purely repulsive, isotropic pair potentials that favor assembly of particles into a wide variety of open 2D lattices. Because this method is based on standard, MD simulation techniques, there is no need to construct or update a large competitor pool during the optimization. Instead, structural information is simply encountered in the simulation and

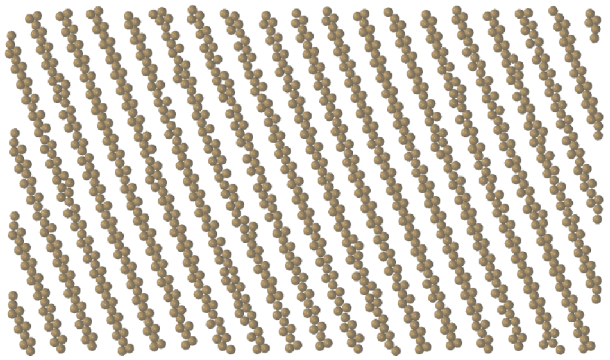

FIG. 5. Self-assembled structure of $u(r)$ discovered with an optimization targeting the TH lattice, but performed without starting each step from a fluid configuration.

utilized in the optimization on-the-fly, via self-assembly, to target a given crystal–including states that would otherwise be hard to predict *a priori* or to incorporate into a competitor pool (such as periodic microphases). In addition to the major simplification from the implicit competitor pool, mechanical stability, and likewise some finite degree of thermal stability, are naturally encoded in the method. However in this approach, it is not known whether the desired lattice is the ground state; rather, the optimization gives insight into what structure will result from self-assembly at nonzero temperature for a given potential. Also, it should be noted that crystals with large kinetic barriers to self-assembly will likely be difficult to treat with this strategy, though it is also reasonable to suspect that such lattices will encounter real-world complications that may preclude their assembly in general–particularly for large nanoparticle to micron-sized colloidal systems.

Intriguing avenues for future work include optimizing other constrained functional forms for $u(r)$, and the use of different types of target structures, such as 3D lattices or binary lattices with multiple interactions. With respect to the former, one strategy to move towards realization of such assemblies would be to constrain $u(r|\boldsymbol{\theta})$ to experimentally motivated models that describe, for example, ligand-coated colloids.[29,30] Some existing interaction models, for micelles[8] for instance, bear resemblance to the simpler potentials presented in this work (HC and SQ in particular). However, simplification of the remaining potentials via greater constraint on the functional form might be necessary to realize the corresponding lattices.

ACKNOWLEDGMENTS

T.M.T. acknowledges support of the Welch Foundation (F-1696) and the National Science Foundation (CBET-1403768). We also acknowledge the Texas Advanced Computing Center (TACC) at the University of Texas

at Austin for providing computing resources used to obtain results presented in this paper.

APPENDIX

A. Derivation of Update Scheme

Here we provide a brief derivation of the update scheme shown in Eqn. 1 of the main text.

The probability of observing a configuration, $\mathbf{R}_i$, in the canonical ensemble is given by the Boltzmann factor normalized by the partition function, or Z. Therefore the probability of observing M statistically independent and identically distributed configurations, $\mathbf{R}_{1:M}$, in the NVT ensemble is given by the product of such terms:

$$P(\mathbf{R}_{1:M}|\boldsymbol{\theta}) \equiv \prod_{i=1}^{M} P(\mathbf{R}_i|\boldsymbol{\theta}) = \prod_{i=1}^{M} \frac{\exp[-\beta U(\mathbf{R}_i|\boldsymbol{\theta})]}{Z(\boldsymbol{\theta})} \quad \text{(A1)}$$

where $U(\mathbf{R}_i|\boldsymbol{\theta})$ is the potential energy for configuration $\mathbf{R}_i$ and $\boldsymbol{\theta}$ is a vector of the tunable values that parametrize the potential. (This quantity is also termed *the likelihood* of $\boldsymbol{\theta}$ given $\mathbf{R}_{1:M}$.)

However, we need to optimize $\boldsymbol{\theta}$ for a given set of configurations sampled from the target simulation (as described in the main text). Therefore, we are actually interested in the quantity $P(\boldsymbol{\theta}|\mathbf{R}_{1:M})$, i.e., the probability of the parameters $\boldsymbol{\theta}$ given the configurations $\mathbf{R}_{1:M}$. Within Bayesian statistics, this quantity is termed the posterior distribution and depends on $P(\mathbf{R}_{1:M}|\boldsymbol{\theta})$ as follows:

$$P(\boldsymbol{\theta}|\mathbf{R}_{1:M}) \equiv \frac{P(\mathbf{R}_{1:M}|\boldsymbol{\theta})P(\boldsymbol{\theta})}{P(\mathbf{R}_{1:M})} \quad \text{(A2)}$$

In Eqn. A2, $P(\boldsymbol{\theta})$ and $P(\mathbf{R}_{1:M})$ are marginalized probabilities, or the probabilities in the absence of any information about the other variable. Therefore, $P(\boldsymbol{\theta})$ has the interpretation of a "prior" distribution–reflecting any prior knowledge pertaining to what the parameters should be–and $P(\boldsymbol{\theta}|\mathbf{R}_{1:M})$ reflects how this prior is updated once observational information ($\mathbf{R}_{1:M}$) is taken into account. The latter is the function that we seek to maximize with respect to $\boldsymbol{\theta}$:

$$\text{argmax}_{\boldsymbol{\theta}} P(\boldsymbol{\theta}|\mathbf{R}_{1:M}) \quad \text{(A3)}$$

In our current work, we assume a uniform distribution for $\boldsymbol{\theta}$. With a flat prior for $\boldsymbol{\theta}$, and because $P(\mathbf{R}_{1:M})$ is $\boldsymbol{\theta}$-independent, it is easy to see that $\text{argmax}_{\boldsymbol{\theta}} P(\boldsymbol{\theta}|\mathbf{R}_{1:M}) = \text{argmax}_{\boldsymbol{\theta}} P(\mathbf{R}_{1:M}|\boldsymbol{\theta})$. Thus, for the remainder of this text we seek to maximize the likelihood $P(\mathbf{R}_{1:M}|\boldsymbol{\theta})$ with respect to $\boldsymbol{\theta}$, and maximization of the posterior distribution under the above assumptions is referred to as the maximum likelihood approach.

In practice, it is easier to maximize the log-likelihood. Taking the natural log of Eqn. A1 and dividing by M yields:

$$\begin{aligned} \frac{1}{M}\ln P(\mathbf{R}_{1:M}|\boldsymbol{\theta}) &\equiv \frac{1}{M}\sum_{i=1}^{M} \ln P(\mathbf{R}_i|\boldsymbol{\theta}) \\ &= -\frac{1}{M}\sum_{i=1}^{M}\Big[\beta U(\mathbf{R}_i|\boldsymbol{\theta})\Big] - \ln Z(\boldsymbol{\theta}) \end{aligned} \quad \text{(A4)}$$

which can be written as

$$\langle \ln P(\mathbf{R}|\boldsymbol{\theta})\rangle_{P_{\text{tgt}}(\mathbf{R})} = -\langle \beta U(\mathbf{R}|\boldsymbol{\theta})\rangle_{P_{\text{tgt}}(\mathbf{R})} - \ln Z(\boldsymbol{\theta}) \quad \text{(A5)}$$

in the large configuration limit, i.e., $M \to \infty$. $P_{\text{tgt}}(\mathbf{R})$ is the probability distribution of the target simulation (from which the configurations are actually sampled).

In order to carry out the log-likelihood maximization, we employ a gradient ascent optimization algorithm,

$$\boldsymbol{\theta}^{(i+1)} = \boldsymbol{\theta}^{(i)} + \alpha\big[\boldsymbol{\nabla}_{\boldsymbol{\theta}}\langle \ln P(\mathbf{R}|\boldsymbol{\theta})\rangle_{P_{\text{tgt}}(\mathbf{R})}\big]_{\boldsymbol{\theta}=\boldsymbol{\theta}^{(i)}} \quad \text{(A6)}$$

where α is the empirically determined step size. From Eqn. A5 and employing the relation $Z(\boldsymbol{\theta}) \equiv \int d\mathbf{R} \exp[-\beta U(\mathbf{R}|\boldsymbol{\theta})]$, we find that

$$\begin{aligned} &\boldsymbol{\nabla}_{\boldsymbol{\theta}}\langle \ln P(\mathbf{R}|\boldsymbol{\theta})\rangle_{P_{\text{tgt}}(\mathbf{R})} = \\ &-\langle \boldsymbol{\nabla}_{\boldsymbol{\theta}}\beta U(\mathbf{R}|\boldsymbol{\theta})\rangle_{P_{\text{tgt}}(\mathbf{R})} + \langle \boldsymbol{\nabla}_{\boldsymbol{\theta}}\beta U(\mathbf{R}|\boldsymbol{\theta})\rangle_{P(\mathbf{R}|\boldsymbol{\theta})} \end{aligned} \quad \text{(A7)}$$

For the specific case of an isotropic pair potential, i.e., $U(\mathbf{R}|\boldsymbol{\theta}) \equiv \frac{1}{2}\sum_{i\neq j}^{N} u(r_{i,j}|\boldsymbol{\theta})$, we can factorize the terms on the righthand side of Eqn. A7 into integrals over the product of the gradient with respect to $\boldsymbol{\theta}$ of the pair potential, $\boldsymbol{\nabla}_{\boldsymbol{\theta}} u(r|\boldsymbol{\theta})$, and an averaged two-point density. Because the potential only depends on pair interactions, only a two-point correlation function is required. Therefore, in two dimensions, the gradient required by Eqn. A6 can be rewritten as

$$\begin{aligned} &\boldsymbol{\nabla}_{\boldsymbol{\theta}}\langle \ln P(\mathbf{R}|\boldsymbol{\theta})\rangle_{P_{\text{tgt}}(\mathbf{R})} = \\ &\pi\rho N \int_0^{\infty} dr r[g(r|\boldsymbol{\theta}) - g_{\text{tgt}}(r)]\boldsymbol{\nabla}_{\boldsymbol{\theta}} u(r|\boldsymbol{\theta}) \end{aligned} \quad \text{(A8)}$$

after employing the definition of the radial distribution function, $g(r)$, performing one of the integrals over all space, and converting to spherical coordinates. Inserting this term into Eqn. A6 and absorbing all constants into α yields Eqn. 1 in the main text.

Finally, note that the lefthand side of Eqn. A5 differs from the Kullback-Leibler divergence in the large sample limit, $\big\langle \ln P_{\text{tgt}}(\mathbf{R})\big\rangle_{P_{\text{tgt}}(\mathbf{R})} - \big\langle \ln P(\mathbf{R}|\boldsymbol{\theta})\big\rangle_{P_{\text{tgt}}(\mathbf{R})}$, only by a constant in $\boldsymbol{\theta}$ and in sign. Therefore minimization of the Kullback-Leibler divergence, a typical starting point for the derivation of relative entropy coarse-graining procedures, is equivalent to Eqn. 1 in the main text.

B. Additional Optimization and Simulation Details

The number of iterations needed in the optimization varied with lattice type, but generally approximately 100-200 iterations were required. Therefore, for the iterative optimization scheme, relatively short simulations of

$8\text{x}10^6$ steps were needed for computational efficiency of the optimization. The first half of these simulations entailed cooling from $T/T^* = 1.5$ to $T/T^* = 1.0$, where T^* is defined as the optimization temperature, yielding a cooling rate of $1.25\text{x}10^{-7}T^*/dt$. (The time step for the MD simulations is also defined in terms of T^*: $dt/\sqrt{\sigma^2 m\beta^*} \approx 0.001$, where $\beta^* = (k_B T^*)^{-1}$.) The second half of the simulation was run at T^*; 667 configurations evenly distributed over the final $1\text{x}10^6$ steps were used to collect statistics to compute $g(r|\boldsymbol{\theta})$.

Subsequent simulations with the optimized potentials were much longer, $1\text{x}10^8$ steps, $5\text{x}10^7$ of which were spent cooling between the melting and freezing temperatures as described in the main text. The cooling rate depended on the gap in temperature between these two end points, varying between $8.3\text{x}10^{-10}T^*/dt$ and $2.5\text{x}10^{-9}T^*/dt$ for the different lattices. We observed no discernible dependence of the self-assembly towards slowing the cooling rate over orders of magnitude. Below the freezing temperature (i.e., after large-scale assembly has already occurred), the annealing schedule was accelerated to $\sim 2\text{x}10^{-8}T^*/dt$ as generally only small-scale local rearrangements occurred in this regime.

	r_n/σ	x/σ	y/σ
HC	1.74	27.7	48.0
SQ	1.42	32.0	32.0
KG	2.24	45.0	26.0
R3	3.11	32.0	96.0
TS	2.42	41.0	41.0
TH	3.35	37.4	64.7

TABLE A1. Cut-offs for $u(r)$ and simulation box sizes for each lattice.

The cut-off for $u(r)$ was systematically chosen to include as few coordination shells as possible while still allowing for the correct lattice to form. The resulting cut-offs are given in Table A1, in addition to the box lengths in the x and y directions for the optimization and subsequent simulations.

If the position of the knots in the Akima spline are defined as the vector $\boldsymbol{r} = \{r_1, r_2, ..., r_n\}$, where r_n is the cut-off for the potential, then each knot has a corresponding tunable parameter θ, i.e., $\boldsymbol{\theta} \equiv \theta(\boldsymbol{r})$. In order to simply enforce the monotonicity constraint, the $\theta(\boldsymbol{r})$ values do not directly encode the knot amplitudes ($\boldsymbol{\kappa}$) but rather indicate the difference between the amplitude of the present knot and the next knot:

$$\kappa(r_i) = \kappa(r_{i+1}) + \theta(r_i) \tag{A9}$$

where $\kappa(r_n)$ is fixed at zero. As a result, enforcing the repulsive constraint is equivalent to requiring that all θ values are non-negative.

The starting guess for the optimization is a power law weighted by step-like smoothing function to minimize the force near and at that cutoff

$$\beta u_0(r) \equiv \frac{1}{2}A\left(\frac{\sigma}{r}\right)^a \left(1 - \tanh\left[k\left(\frac{r}{\sigma} - \frac{r_s}{\sigma}\right)\right]\right) \tag{A10}$$

where A is the dimensionless amplitude, k controls the steepness of the switching function, a is the power, r_s is the center of the smoothing function, and σ is the nearest neighbor crystal distance. In this publication we set $A \equiv 1.8$, $k \equiv 8.9$, $a \equiv 5$ and r_s is set near to, but before, the potential cut-off. However, we have not found the optimization to be sensitive to the details of the initial guess. Moreover, Ref. 20 in the main text notes that for linear parameters, such as $\boldsymbol{\theta}$ as defined above, there is only one maximum on the landscape.

REFERENCES

[1] B. Chacko, C. Chalmers, and A. J. Archer, "Two-dimensional colloidal fluids exhibiting pattern formation," J. Chem. Phys. **143**, 244904 (2015).

[2] Y. Zhuang and P. Charbonneau, "Equilibrium phase behavior of the square-well linear microphase-forming model," J. Phys. Chem. B **120**, 6178–6188 (2016).

[3] B. A. Lindquist, R. B. Jadrich, and T. M. Truskett, "Assembly of nothing: equilibrium fluids with designed structured porosity," Soft Matter **12**, 2663–2667 (2016).

[4] G. Zhang, F. H. Stillinger, and S. Torquato, "Probing the limitations of isotropic pair potentials to produce ground-state structural extremes via inverse statistical mechanics," Phys. Rev. E **88**, 042309 (2013).

[5] A. Jain, J. R. Errington, and T. M. Truskett, "Dimensionality and design of isotropic interactions that stabilize honeycomb, square, simple cubic, and diamond lattices," Phys. Rev. X **4**, 031049 (2014).

[6] E. Marcotte, F. H. Stillinger, and S. Torquato, "Optimized monotonic convex pair potentials stabilize low-coordinated crystals," Soft Matter **7**, 2332–2335 (2011).

[7] W. D. Piñeros, M. Baldea, and T. M. Truskett, "Breadth versus depth: Interactions that stabilize particle assemblies to changes in density or temperature," J. Chem. Phys. **144**, 084502 (2016).

[8] T. Dotera, T. Oshiro, and P. Ziherl, "Mosaic two-lengthscale quasicrystals," Nature **506**, 208–211 (2014).

[9] K. Barkan, H. Diamant, and R. Lifshitz, "Stability of quasicrystals composed of soft isotropic particles," Phys. Rev. B **83**, 172201 (2011).

[10] S. Torquato, "Inverse optimization techniques for targeted self-assembly," Soft Matter **5**, 1157–1173 (2009).

[11] A. Jain, J. A. Bollinger, and T. M. Truskett, "Inverse methods for material design," AIChE J. **60**, 2732–2740 (2014).

[12] A. Jain, J. R. Errington, and T. M. Truskett, "Inverse design of simple pairwise interactions with low-coordinated 3D lattice ground states," Soft Matter **9**, 3866–3870 (2013).

[13] E. Marcotte, F. H. Stillinger, and S. Torquato, "Communication: Designed diamond ground state via optimized isotropic monotonic pair potentials," J. Chem. Phys. **138**, 061101 (2013).

[14] W. Humphrey, A. Dalke, and K. Schulten, "VMD – Visual Molecular Dynamics," J. Mol. Graphics **14**, 33–38 (1996).

[15] M. Z. Miskin, G. Khaira, J. J. de Pablo, and H. M. Jaeger, "Turning statistical physics models into materials design engines," Proc. Nat. Acad. Sci. **113**, 34–39 (2016).

[16] R. L. Marson, T. D. Nguyen, and S. C. Glotzer, "Rational design of nanomaterials from assembly and reconfigurability of polymer-tethered nanoparticles," MRS Communications **5**, 397–406 (2015).

[17]A. F. Hannon, Y. Ding, W. Bai, C. A. Ross, and A. Alexander-Katz, "Optimizing topographical templates for directed self-assembly of block copolymers via inverse design simulations," Nano Lett. **14**, 318–325 (2014).
[18]A. W. Long and A. L. Ferguson, "Nonlinear machine learning of patchy colloid self-assembly pathways and mechanisms," J. Phys. Chem. B **118**, 4228–4244 (2014).
[19]D. Barber, *Bayesian Reasoning and Machine Learning* (Cambridge University Press, 2012).
[20]M. S. Shell, "The relative entropy is fundamental to multiscale and inverse thermodynamic problems," J. Chem. Phys. **129**, 144108 (2008).
[21]W. G. Noid, "Perspective: Coarse-grained models for biomolecular systems," J. Chem. Phys. **139**, 090901 (2013).
[22]A. Chaimovich and M. S. Shell, "Coarse-graining errors and numerical optimization using a relative entropy framework," J. Chem. Phys. **134**, 094112 (2011).
[23]M. C. Rechtsman, F. H. Stillinger, and S. Torquato, "Optimized interactions for targeted self-assembly: Application to a honeycomb lattice," Phys. Rev. Lett. **95**, 228301 (2005).
[24]$\alpha \approx 0.02$ worked well for the cases studied here.
[25]The magnitude of the spring constant is chosen such that the radial distribution function has sharp crystalline features, but the peaks are still integrable. Generally this corresponded to a range of 1600-2800 k_BT/σ^2, but we do not anticipate that the exact value is critical as long as the $g(r)$ is consistent with a stable crystal and not a fluid.
[26]D. Van Der Spoel, E. Lindahl, B. Hess, G. Groenhof, A. E. Mark, and H. J. C. Berendsen, "Gromacs: Fast, flexible, and free," J. Comput. Chem. **26**, 1701–1718 (2005).
[27]S. Plimpton, "Fast parallel algorithms for short-range molecular dynamics," J. Comput. Phys. **117**, 1–19 (1995).
[28]A Nosé-Hoover thermostat with a time constant of $\tau = 100dt$ is employed, where dt is the time step.
[29]J. Mewis and N. J. Wagner, *Colloidal Suspension Rheology* (Cambridge University Press, 2013).
[30]C. N. Likos, K. A. Vaynberg, H. Löwen, and N. J. Wagner, "Colloidal stabilization by adsorbed gelatin," Langmuir **16**, 4100–4108 (2000).